# Electronic Structure Calculations of Quantum Dots Consisting of Tens of Millions of Atoms on High Performance Beowulf Clusters

Fabiano Oyafuso, Gerhard Klimeck, R. Chris Bowen, Timothy B. Boykin[#]
*Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA 91109*
*[#]University of Alabama at Huntsville, Dept. of Electrical Engineering, AL, 35899*

Non-atomistic models such as **k.p** are often used to characterize the electronic structures of quantum dots. However, such jellium-like models are fundamentally not well suited for the atomistic representation of nano-scale features such as interfaces and disorder. A realistic atomic simulation, however, is computationally taxing, since it must include millions of atoms. For example, the modeling of a horizontal array of four self-assembled InAs quantum dots of 30 nm diameter and 5 nm height embedded in GaAs of buffer width 5 nm and separated by 20 nm requires a simulation domain encompassing 5.2 million atoms. The large problem size of these applications, then, necessitates computation on large parallel computers.

In this work, we describe an extension of the NEMO 1-D device simulator developed at Raytheon/TI to 3D in order to model quantum dot structures on high performance commodity clusters (Beowulfs). Our simulation employs a nearest-neighbor tight-binding model with a 20 orbital basis, consisting of s, p, and d orbitals, associated with each atomic lattice site. The coupling energies between these orbitals are computed by using a genetic algorithm package to determine a best fit to experimentally measurable bulk properties such as bandgaps, effective masses and strain-induced shifts. This coupling, represented by a 20×20 matrix, is dependent on the bond length and, in strained systems, is different for each nearest neighbor cation-anion pair.

Computation of the electronic structure of a typical system involves the diagonalization of a sparse Hamiltonian of order $\sim 10^8$ though a customized parallel Lanczos solver. Data are partitioned such that each processor holds information relevant only to the set of atoms associated with that processor. Furthermore, atoms are mapped to processors in such a way as to preserve their geometric connectivity, so that our nearest-neighbor model results only in nearest-neighbor communication among processors. This scheme allows for a simple, efficient 1D chain network topology in which all communication is synchronized right to left, then left to right. Data structures specific to zincblende are used to minimize storage requirements by eliminating redundancy due to symmetry, so that on a 64 processor system with 1 GB/processor, nanostructures consisting of up to 15 million atoms can be modeled before there is significant swapping. Even larger systems can be simulated by recomputing the Hamiltonian on the fly, although at a cost of a factor of 2 to 4 in execution time depending on platform. Scaling results illustrating performance as a function of number of processors and problem size will be presented.

Roughly a third of the CPU time is spent on computing the mechanical strain on which the electronic calculation depends. The interatomic distances are determined by minimizing the mechanical strain within a nearest neighbor valence force model using a parallelized conjugate gradient based method. Both fixed and periodic boundary conditions are allowed, and in the case of periodic boundary conditions the strain energy is also minimized with respect to the overall periodicity. The resulting lattice constant of $In_xGa_{1-x}As$ alloys shows good agreement with the VCA lattice constant.

An important advantage of the atomistic model is in the characterization of disorder. We will present results showing that even in the absence of carrier-carrier or carrier-phonon scattering or dot size distributions, a spread in transition energy of several meV can still be expected simply due to the variations in the random distribution of cations. Our work also demonstrates that simple hard-wall boundary conditions can strongly affect the converged eigenenergies if the surrounding buffer is sufficiently thin so that care must be taken to ensure proper convergence.